

Fundus2Video: Cross-Modal Angiography Video Generation from Static Fundus Photography with Clinical Knowledge Guidance

Weiye Zhang¹ , Siyu Huang², Jiancheng Yang³, Ruoyu Chen¹, Zongyuan Ge⁴, Yingfeng Zheng⁵, Danli Shi¹ , and Mingguang He¹

¹ The Hong Kong Polytechnic University, Kowloon, Hong Kong

² Clemson University, South Carolina, USA

³ École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland

⁴ Monash University, Melbourne, Australia

⁵ Sun Yat-sen University, Guangzhou, China
danli.shi@polyu.edu.hk

Abstract. Fundus Fluorescein Angiography (FFA) is a critical tool for assessing retinal vascular dynamics and aiding in the diagnosis of eye diseases. However, its invasive nature and less accessibility compared to Color Fundus (CF) images pose significant challenges. Current CF to FFA translation methods are limited to static generation. In this work, we pioneer dynamic FFA video generation from static CF images. We introduce an autoregressive GAN for smooth, memory-saving frame-by-frame FFA synthesis. To enhance the focus on dynamic lesion changes in FFA regions, we design a knowledge mask based on clinical experience. Leveraging this mask, our approach integrates innovative knowledge mask-guided techniques, including knowledge-boosted attention, knowledge-aware discriminators, and mask-enhanced patch-NCE loss, aimed at refining generation in critical areas and addressing the pixel misalignment challenge. Our method achieves the best FVD of 1503.21 and PSNR of 11.81 compared to other common video generation approaches. Human assessment by an ophthalmologist confirms its high generation quality. Notably, our knowledge mask surpasses supervised lesion segmentation masks, offering a promising non-invasive alternative to traditional FFA for research and clinical applications. The code is available at <https://github.com/Michi-3000/Fundus2Video>.

Keywords: Video Generation · Generative Adversarial Network · Autoregressive Generation · Retinal Fundus Photography · Fluorescence Angiography.

1 Introduction

Fundus Fluorescein Angiography (FFA) is an essential examination in ophthalmology clinics, providing a dynamic view of retinal blood flow and lesion changes. It offers critical insights into retinal circulatory dynamics, aiding in the identification of conditions such as diabetic retinopathy, hypertensive retinopathy, and

macular degeneration [13]. Unlike Color Fundus (CF) images, FFA videos capture the dynamic filling process and real-time changes in retinal vascular abnormalities with greater clarity and depth, thereby enhancing diagnostic precision and facilitating a deeper understanding of disease progression and treatment response. However, due to its invasive nature and potential side effects, FFA’s use is limited for certain individuals. In contrast, CF photography is non-invasive, readily available [27], and has been utilized in some deep-learning methods [6,21] for disease diagnosis. Therefore, generating realistic FFA videos from CF images holds significant research and application potential.

When considering the generative models for FFA synthesis, the majority of existing methods [23,11,12,16,20] focus on specific phases, like the venous and late phase, using various Generative Adversarial Networks (GAN). However, they overlook the changes occurring throughout the entire FFA process, which includes multiple phases. While some approaches [2] can generate multiple discrete FFA images from different phases simultaneously, they still cannot capture the fully dynamic changes of retinal structures and lesions. Capturing lesional changes accurately is another challenge in FFA generation. While using lesion labels for conditional supervision could potentially enhance image details, the manual annotation of these labels is highly time-consuming and impractical for segmenting all possible lesion changes. Additionally, the time-consuming nature of FFA procedures makes it difficult to align FFA images precisely with CF images in clinical practice, due to blinking and movement, even with good patient cooperation [4,7]. This misalignment poses a significant challenge for pixel-to-pixel-based video generation processes.

To tackle these challenges, we propose a model leveraging an image-to-image GAN framework, specifically pix2pixHD [25], to generate smooth and stable FFA videos from single CF images autoregressively. Through clinical knowledge analysis of ground-truth FFA series, regions with significant lesion changes during the early and late FFA series examination are lesional changes, reflecting the damage in vascular or retinal pigment epithelium structure [3,28]. The larger the changes, the more important they are. Leveraging this insight, we design a knowledge mask that requires no additional manual labeling and enhances the generation of regions with high variability. Using this mask, we introduce novel knowledge mask-guided techniques into the baseline model to guide the model to focus more on key regions during learning and generation. Specifically, we propose a mask-enhanced patchNCE loss to address the pixel misalignment issue. This model holds the potential to generate FFA videos from CF images to other modalities and improve downstream tasks [1,22,19].

In summary, our research contributes as follows: **1.** We are the first to generate dynamic FFA video directly from CF images, marking a significant advancement in ophthalmic imaging. Specifically, we introduce Fundus2Video, an autoregressive GAN architecture tailored for frame-by-frame FFA video synthesis from CF images. This architecture optimizes memory usage and ensures smooth output. **2.** We introduce a knowledge mask derived from clinical insights to enhance focus on regions undergoing significant changes during dynamic FFA processes.




Fig. 1. Proposed Fundus2Video. (a) The overall architecture. Generator G generates one frame at a time, taking the output from the previous time step and the CF image as input. During the training phase, unsupervised knowledge masks guide the entire network. (b) The design of the mask-enhanced patchNCE loss.

This eliminates manual labeling and improves generation in areas like lesions and blood vessels. **3.** With this mask, we implement knowledge mask (KM)-guided techniques. We introduce knowledge-mask-guided attention and knowledge-aware discriminators for specific supervision on regions of lesion regions. To address the pixel misalignment challenge between CF images and ground-truth FFA series in critical areas, we employ a newly designed mask-enhanced patchNCE loss.

2 Methods

2.1 Overview

We aim to generate a realistic FFA video $\hat{Y}$ from a given CF image x , with the ground-truth FFA video during training represented as Y . Considering the temporal nature of FFA series, we adopt an autoregressive GAN architecture to capture temporal dependencies and generate coherent video sequences. An autoregressive GAN generates image samples sequentially, conditioning each new image on previously generated images and additional inputs. In our context of generating FFA videos from CF images, our autoregressive GAN, named Fundus2Video, based on the image-to-image translation GAN pix2pixHD, sequentially generates each frame of the FFA video, incorporating the CF image itself and the preceding frames. Building upon the generator, discriminator, and loss designs of pix2pixHD, our approach incorporates specific modifications to enable autoregressive and smooth generation. The architecture is as shown in Fig. 1 (a).

To ensure smooth output in Fundus2Video, we incorporate multi-frame input and smoothing techniques for longer temporal considerations. Specifically, we input three consecutive frames from the ground-truth FFA series to the model in




Fig. 2. The definition of the knowledge mask. Left: The unsupervised process of obtaining the mask. The knowledge mask covers the same pathological areas as the expert-labeled mask. Right: Generated results with and without the knowledge mask.

a sliding window fashion to provide longer temporal context for each generated frame. Instead of generating each frame independently, we aggregate the generated frames over a sliding window and perform triple-frame averaging. This approach smooths out abrupt transitions between frames and ensures continuity in the generated video sequence.

2.2 Unsupervised Clinically Supported Knowledge Mask

The baseline Fundus2Video can generate smooth and continuous FFA videos. However, it falls short of accurately depicting details like lesions and critical structures as shown in Fig. 2 right, which are of utmost clinical importance. To address this, we leveraged clinical insights to analyze ground-truth FFA videos, which tell us regions undergoing significant morphological changes during the FFA process often corresponded to crucial lesions or retinal structure areas that pose challenges for the model. The theoretical basis is from [3]:

- *During the FFA process, as the fluorescent dye flows through retinal vessels, significant leakage always occurs around the lesions, leading to visible differences between early and late stages.*

Building upon this knowledge, we devised a simple binary mask by computing the difference between the first frame (representing the arterial phase) and the last frame (representing the late phase) and setting a specific threshold δ determined through comparative experiments, which can be formulated as $m = \delta(Y_0 - Y_T)$, where Y_0 represents the first frame of the ground-truth FFA video and Y_T represents the last frame. The process is depicted in Fig. 2. Unlike supervised lesion/structure segmentation masks, this knowledge mask requires no additional manual annotation or segmentation model training. It can be easily derived from raw data, making it simple yet effective.

2.3 Knowledge Mask-Guided Video Generation

Knowledge-boosted Attention. Some types of lesions may be challenging to detect in CF images due to low contrast, leading to synthesized FFA slices lacking

details in these areas. To address this limitation and improve the generator’s ability to capture specific regions, we introduce additional supervision into the learning process. Our approach, termed knowledge-boosted attention, involves guiding the network’s attention toward focal regions during training. To quantify this guidance, we define an attention loss $\mathcal{L}_{Att}$ as follows:

$$\mathcal{L}_{Att}(A, m) = \frac{1}{n} \sum_i (A^i - m^i)^2. \quad (1)$$

Here, m represents the knowledge mask described in Section 2.2. A denotes the attention map obtained by element-wise multiplication of the semantic-rich activation map f_l from the last convolutional layer l in the generator and the mask m . We then apply a rectified linear operation to A , resulting in $A = ReLU(f_l \odot m)$.

Mask-enhanced PatchNCE Losses. To address pixel misalignment in ground-truth FFA series and CF images caused by motion artifacts during acquisition, we introduce the PatchNCE loss [14], inspired by contrastive learning techniques known for boosting model robustness against label noise. However, we observed that the model’s primary focus should be on reducing jitter in clinically relevant regions, such as lesions and vasculature, which are of greater clinical significance. To further tackle this issue, we propose mask-enhanced PatchNCE losses as a replacement for traditional PatchNCE losses. This method extends traditional PatchNCE losses by incorporating a knowledge mask m , highlighting critical regions within the FFA series. Mathematically, the proposed mask-enhanced PatchNCE losses are based on the InfoNCE loss, which is defined as:

$$\mathcal{L}_{InfoNCE}(v, v^+, v^-) = -\log \left(\frac{e^{sim(v, v^+)}}{e^{sim(v, v^+)} + \sum_{j=1}^N e^{sim(v, v_j^-)}} \right). \quad (2)$$

Here, v , v^+ , and v^- represent the embeddings of the anchor, positive, and negative samples, respectively.

The mask-enhanced unsupervised PatchNCE (UP) loss compares the anchor patch $z_{\hat{Y}}$ in the generated output with a corresponding positive patch z_X from the input CF image and negative patches z_X^- , under the guidance of knowledge mask m . It is defined as:

$$L_{Masked\ UP} = L_{InfoNCE}(m \odot z_{\hat{Y}}, m \odot z_X, m \odot z_X^-), \quad (3)$$

where $\odot$ denotes element-wise multiplication. In contrast, the mask-enhanced supervised PatchNCE (SP) loss ensures consistency between generated and ground-truth patches. It designates the corresponding patch in the ground-truth image z_Y as positive, while non-corresponding patches z_Y^- are considered negatives. It’s defined as:

$$L_{Masked\ SP} = L_{InfoNCE}(m \odot z_{\hat{Y}}, m \odot z_Y, m \odot z_Y^-). \quad (4)$$

The illumination is shown in Fig. 1 (b). By integrating the knowledge mask into the PatchNCE loss, our method directs the model’s focus during training, improving its ability to capture clinically significant features.

Knowledge-aware Discriminators. We employ 3 discriminators $D = \{D_1, D_2, D_3\}$ [9,25] with the same patchGAN architecture [10] to evaluate images at scales of 1, 0.5 and 0.25 for different receptive fields. The discriminator objective function for D_k with generator G is given by:

$$\mathcal{L}_{D_k}(a, b, G(a)) = E_{a,b}[\log D_k(a, b)] + E_a[\log(1 - D_k(a, G(a)))], \quad (5)$$

where a , b and $G(a)$ are the input, ground-truth, and generated images.

However, solely discriminating the entire image may not ensure the authenticity of lesion regions in generated FFA frames. Hence, we introduce discrimination guided by knowledge across scales. By combining knowledge masks m with corresponding FFA images, we tailor inputs for the discriminators to focus on lesions. According to Eq. 5, the combined discriminator loss $\mathcal{L}_{GAN}(G, D_k)$ for scale k is defined as:

$$\mathcal{L}_{GAN}(G, D_k) = \mathcal{L}_{D_k}(x, y, G(x)) + \mathcal{L}_{D_k}(x \odot m, y \odot m, G(x) \odot m), \quad (6)$$

where $\odot$ denotes element-wise multiplication, and x and y are the input CF images and ground-truth FFA images, respectively.

Consequently, the final loss function is as follows:

$$\mathcal{L} = \lambda_{UP} \mathcal{L}_{MaskedUP} + \lambda_{SP} \mathcal{L}_{MaskedSP} + \lambda_{Att} \mathcal{L}_{Att} + \lambda_{GAN} \mathcal{L}_{GAN}. \quad (7)$$

3 Experiments

Dataset. Our dataset comprises 350 CF images and 18,180 corresponding FFA images from 350 anonymous patients sampled from a large paired dataset. The FFA images were obtained using Zeiss FF450 Plus and Heidelberg Spectralis systems, with a resolution of 768×768 pixels. Meanwhile, the CF images were captured by Topcon TRC-50XF and Zeiss FF450 Plus instruments, with resolutions ranging from 1,110×1,467 to 2,600×3,200 pixels. The Institutional Review Board approved the study.

Implementation Details. The final objective function (Eq. 7) was utilized to train the generative model, with λ_{UP} , λ_{SP} , λ_{Att} , and λ_{GAN} set to 1, 1, 4, and 2, respectively. The threshold δ for obtaining the knowledge mask was set to 45. During training, each ground-truth FFA series produced 12 frames, with 4 slices randomly selected from the vascular, venous, and late phases, respectively. Data augmentation techniques including random cropping, scaling, and color augmentation. The input images were resized to 512×512. Additionally, the model was trained to randomly select either generated or ground-truth frames as input, enhancing its adaptability and robustness. We employed the Adam optimizer with




Fig. 3. Qualitative comparison of the methods. Frames are sampled from the 12-frame video. Areas in red boxes denote significant lesions. It can be observed that the KM-guided Fundus2Video exhibits best performance in generating critical lesions.

$\beta_1 = 0.5$ and $\beta_2 = 0.999$, adjusting the learning rate every 50 iterations using the PyTorch [17] lr scheduler. The initial learning rate was set to $2e-3$, with a batch size of 1. Training was conducted for 50 epochs on an NVIDIA GeForce RTX3090. For evaluation, 70% of the data was reserved for training at the patient level, while the remaining data was evenly split into validation and test sets.

Evaluation Criteria. Our video evaluation criteria include Fréchet Video Distance (FVD) [24], Structural Similarity Index (SSIM) [26], Peak Signal-to-Noise Ratio (PSNR) [8], and Learned Perceptual Image Patch Similarity (LPIPS) [29]. They measure feature distribution similarity, video structural similarity, reconstruction quality, and perceptual similarity, respectively.

Model Comparisons. We evaluate Fundus2Video against existing image-to-video translation methods, including the auto-encoder-based Seg2vid [15], and

Table 1. Comparison of the methods. M stands for masks.

Models	Mask Type	Proposed Techniques				FVD↓	SSIM↑	PSNR↑	LPIPS↓
		PatchNCE loss L_P	Mask-enhanced PatchNCE loss $L_{Masked\ P}$	Knowledge -boost attention $\mathcal{L}_{Att}$	Knowledge -aware discriminators				
Seg2Vid[15]	-	-	-	-	-	2302.15	0.2930	10.23	0.2451
Med-ddpm[5]	-	-	-	-	-	2410.54	0.2305	10.59	0.2513
ConsistI2V[18]	-	-	-	-	-	2108.33	0.2662	10.71	0.2498
Fundus2Video	Knowledge M	✗	✗	✗	✗	1804.25	0.3225	11.11	0.2213
	Knowledge M	✓	✗	✗	✗	1611.21	0.3625	11.41	0.2162
	Knowledge M	✗	✓	✗	✗	1527.94	0.3738	11.76	0.2093
	Knowledge M	✗	✗	✓	✗	1701.30	0.3694	11.20	0.2133
	Knowledge M	✗	✗	✗	✓	1664.42	0.3442	11.36	0.2166
	GT Lesion Seg M	✓	✓	✓	✓	1586.35	0.3688	11.23	0.2136
	Knowledge M	✓	✓	✓	✓	1503.21	0.3814	11.81	0.2001

the diffusion model-based Med-ddpm [5] and ConsistI2V [18]. Table 1 shows our model’s superior performance across all metrics. Qualitative comparison in Fig. 3 reveals clearer images and discernible lesion areas in our approach versus others.

Ablation Studies. Additionally, we conduct comprehensive ablation studies to assess the effectiveness of our proposed knowledge mask and related techniques, detailed in the latter part of Table 1 and Fig. 3. Firstly, we show that our designed mask-enhanced patchNCE loss, knowledge-boost attention, and knowledge-aware discriminators, when combined with mask information, outperform the baseline Fundus2Video. Moreover, our mask-enhanced patchNCE loss yields better results than patchNCE loss alone. Secondly, by replacing the knowledge mask with the ground-truth lesion segmentation mask for comparison, we observe that utilizing our KM-guided techniques can enhance performance even with the lesion segmentation mask. However, our knowledge mask yields better results without the need for additional training or labeling efforts.

Human Assessment. An ophthalmologist reviewed the results of all methods in Table 1 and found that our Fundus2Video significantly outperformed others. Then the ophthalmologist conducted a quality assessment of 50 randomly selected FFA videos generated by Fundus2Video from the test set, evaluating them based on their corresponding CF images and ground-truth FFA videos. The evaluation focused on vascular perfusion, lesion dynamics, overall coherence, stability, and presence of artifacts. Scores ranged from 1 to 5, with 1 indicating excellent quality and 5 indicating very poor quality. Our model received a score of 2.12 with a standard deviation of 1.07, indicating good overall quality of the generated videos.

4 Conclusion

In this study, we propose Fundus2Video, which pioneers dynamic FFA video generation from static CF images using an autoregressive GAN architecture. With a

knowledge mask derived from clinical experience, we enhance focus on dynamic lesion regions, outperforming supervised lesion segmentation masks. Our method incorporates knowledge-boosted attention, knowledge-aware discriminators, and mask-enhanced patchNCE loss to address challenges in lesion generation and pixel misalignment. Fundus2Video emerges as a promising alternative to traditional FFA, surpassing recent state-of-the-art approaches with its non-invasive, intuitive, and dynamic features.

Acknowledgments. The study was supported by the Global STEM Professorship Scheme (P0046113) and the Start-up Fund for RAPs under the Strategic Hiring Scheme (P0048623) from HKSAR. The sponsors or funding organizations had no role in the design or conduct of this research.

Disclosure of Interests. A patent has been filed for this innovation (CN 202410360491.4).

References

1. Chen, R., Zhang, W., Song, F., Yu, H., Cao, D., Zheng, Y., He, M., Shi, D.: Translating color fundus photography to indocyanine green angiography using deep-learning for age-related macular degeneration screening. *npj Digital Medicine* **7**(1), 34 (2024)
2. Chen, Y., He, Y., Li, W., Wang, J., Li, P., Xing, L., Zhang, X., Shi, G.: Series-parallel generative adversarial network architecture for translating from fundus structure image to fluorescence angiography. *Applied Sciences* **12**(20), 10673 (2022)
3. Comin, C.H., Tsiurukis, D.I., Sun, Y., Xu, X.: Quantification of retinal blood leakage in fundus fluorescein angiography in a retinal angiogenesis model. *Scientific Reports* **11**(1), 19903 (2021)
4. De Carlo, T.E., Romano, A., Waheed, N.K., Duker, J.S.: A review of optical coherence tomography angiography (octa). *International journal of retina and vitreous* **1**, 1–15 (2015)
5. Dorjsembe, Z., Pao, H.K., Odonchimed, S., Xiao, F.: Conditional diffusion models for semantic 3d medical image synthesis. *arXiv preprint arXiv:2305.18453* (2023)
6. Faust, O., Acharya U, R., Ng, E.Y.K., Ng, K.H., Suri, J.S.: Algorithms for the automated detection of diabetic retinopathy using digital fundus images: a review. *Journal of medical systems* **36**, 145–157 (2012)
7. Freeman, W.R., Bartsch, D.U., Mueller, A.J., Banker, A.S., Weinreb, R.N.: Simultaneous indocyanine green and fluorescein angiography using a confocal scanning laser ophthalmoscope. *Archives of Ophthalmology* **116**(4), 455–463 (1998)
8. Huynh-Thu, Q., Ghanbari, M.: Scope of validity of psnr in image/video quality assessment. *Electronics letters* **44**(13), 800–801 (2008)
9. Iizuka, S., Simo-Serra, E., Ishikawa, H.: Globally and locally consistent image completion. *ACM Transactions on Graphics (ToG)* **36**(4), 1–14 (2017)
10. Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1125–1134 (2017)
11. Kamran, S.A., Hossain, K.F., Tavakkoli, A., Zuckerbrod, S.L.: Attention2angiogan: Synthesizing fluorescein angiography from retinal fundus images using generative adversarial networks. In: *2020 25th International Conference on Pattern Recognition (ICPR)*. pp. 9122–9129. IEEE (2021)

12. Kamran, S.A., Hossain, K.F., Tavakkoli, A., Zuckerbrod, S.L., Sanders, K.M., Baker, S.A.: Rv-gan: Segmenting retinal vascular structure in fundus photographs using a novel multi-scale generative adversarial network. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VIII 24. pp. 34–44. Springer (2021)
13. Kylstra, J.A., Brown, J.C., Jaffe, G.J., Cox, T.A., Gallemore, R., Greven, C.M., Hall, J.G., Eifrig, D.E.: The importance of fluorescein angiography in planning laser treatment of diabetic macular edema. *Ophthalmology* **106**(11), 2068–2073 (1999)
14. Li, F., Hu, Z., Chen, W., Kak, A.: Adaptive supervised patchnce loss for learning h&e-to-ihc stain translation with inconsistent groundtruth image pairs. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2023. pp. 632–641. Springer Nature Switzerland, Cham (2023)
15. Pan, J., Wang, C., Jia, X., Shao, J., Sheng, L., Yan, J., Wang, X.: Video generation from single semantic label map. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3733–3742 (2019)
16. Park, K.B., Choi, S.H., Lee, J.Y.: M-gan: Retinal blood vessel segmentation by balancing losses through stacked deep fully convolutional networks. *IEEE Access* **8**, 146308–146322 (2020)
17. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimselshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
18. Ren, W., Yang, H., Zhang, G., Wei, C., Du, X., Huang, S., Chen, W.: Consisti2v: Enhancing visual consistency for image-to-video generation. *arXiv preprint arXiv:2402.04324* (2024)
19. Shi, D., He, S., Yang, J., Zheng, Y., He, M.: One-shot retinal artery and vein segmentation via cross-modality pretraining. *Ophthalmology Science* **4**(2), 100363 (2024)
20. Shi, D., Zhang, W., He, S., Chen, Y., Song, F., Liu, S., Wang, R., Zheng, Y., He, M.: Translation of color fundus photography into fluorescein angiography using deep learning for enhanced diabetic retinopathy screening. *Ophthalmology Science* **3**(4), 100401 (2023)
21. Sinthanayothin, C., Boyce, J.F., Williamson, T.H., Cook, H.L., Mensah, E., Lal, S., Usher, D.: Automated detection of diabetic retinopathy on digital fundus images. *Diabetic medicine* **19**(2), 105–112 (2002)
22. Song, F., Zhang, W., Zheng, Y., Shi, D., He, M.: A deep learning model for generating fundus autofluorescence images from color fundus photography. *Advances in ophthalmology practice and research* **3**(4), 192–198 (2023)
23. Tavakkoli, A., Kamran, S.A., Hossain, K.F., Zuckerbrod, S.L.: A novel deep learning conditional generative adversarial network for producing angiography images from retinal fundus photographs. *Scientific Reports* **10**(1), 1–15 (2020)
24. Unterthiner, T., van Steenkiste, S., Kurach, K., Marinier, R., Michalski, M., Gelly, S.: Fvd: A new metric for video generation (2019)
25. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8798–8807 (2018)

26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
27. Yannuzzi, L.A., Ober, M.D., Slakter, J.S., Spaide, R.F., Fisher, Y.L., Flower, R.W., Rosen, R.: Ophthalmic fundus imaging: today and beyond. *American journal of ophthalmology* **137**(3), 511–524 (2004)
28. Yannuzzi, L.A., Rohrer, K.T., Tindel, L.J., Sobel, R.S., Costanza, M.A., Shields, W., Zang, E.: Fluorescein angiography complication survey. *Ophthalmology* **93**(5), 611–617 (1986)
29. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)